

(12) UK Patent Application (19) GB (11) 2 366 050 (13) A

(43) Date of A Publication 27.02.2002

(21) Application No 0108635.4

(22) Date of Filing 05.04.2001

(30) Priority Data

(31) 09547402

(32) 11.04.2000

(33) US

(71) Applicant(s)

Hewlett-Packard Company
(Incorporated in USA - Delaware)
3000 Hanover Street, Palo Alto, California 94304,
United States of America

(72) Inventor(s)

Hugh F Mahon
Fredrick Roeling

(74) Agent and/or Address for Service

Carpmaels & Ransford
43 Bloomsbury Square, LONDON, WC1A 2RA,
United Kingdom

(51) INT CL⁷

G06F 11/34

(52) UK CL (Edition T)

G4A AFMP

(56) Documents Cited

EP 0457110 A2

WO 01/20456 A1

WO 00/47003 A1

(58) Field of Search

UK CL (Edition S) G4A AFMD AFMP

INT CL⁷ G06F 11/30 11/32 11/34

ONLINE: WPI, EPODOC, PAJ, INSPEC, IBM TDB

(54) Abstract Title

Aggregation of log data from different operating systems into a central data log

(57) Apparatus and methods for maintaining a centralized data log [105] for a distributed computer system [100] utilizing more than one type of operating system [125,135,165]. The present patent document discloses techniques for transferring and storing log data [300] across different platforms and the aggregation of that log data [300] into one location wherein the processes detecting the log data [300] are executed by operating systems [125,135,165] which are not limited to being of the same type. Thus, this aggregation mechanism is designed to allow multiple processes [130,160] operating on diverse kinds of systems [140,170] to log to a central system [115], which itself may be on any kind of system. An administrator can monitor from a single source the operation of a distributed computer system [100], as for example a distributed management tool, whose components may be distributed across a network [190] and operating on multiple, geographically dispersed computers [140,170].




FIG. 1




FIG. 1




FIG. 2




FIG. 3

4/4

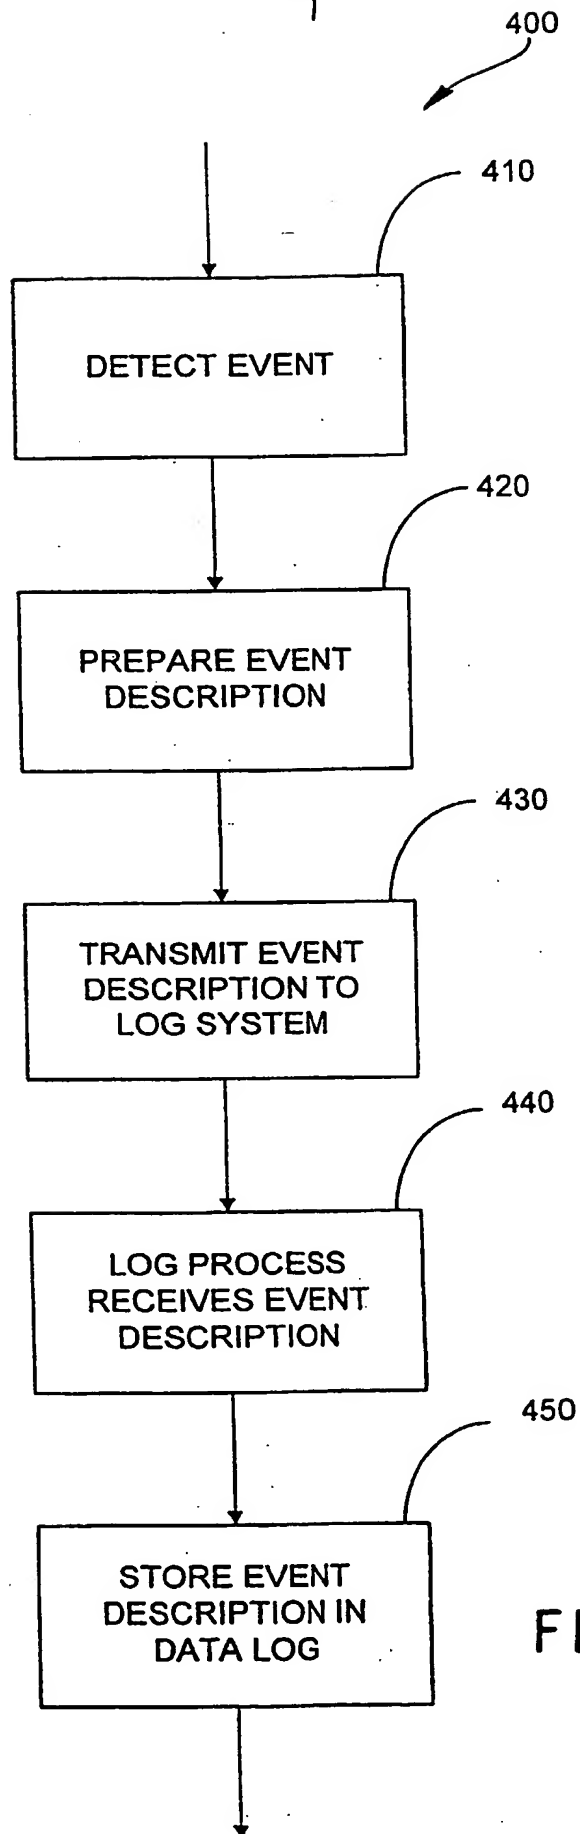


FIG. 4

AGGREGATION OF LOG DATA INTO CENTRALIZED DATA LOG

5

FIELD OF THE INVENTION

The present invention relates generally to networks of computer systems and, more particularly, to the logging of information regarding the activities of the system.

10

BACKGROUND OF THE INVENTION

In order to monitor the progress of any system, it is desirable to have information about the activities of the system. Such information, provided in a manner which collects it over time into one location is called a log. The advent of distributed systems operating on networks, in particular the Internet, has presented new difficulties due to the fact that each individual system maintains its own separate, local log. Thus, in order to investigate operation of the system as a whole, system administrators have been forced to open and read several different logs. The system administrator's job is made especially difficult in trying to correlate the timing of several events which were recorded in differing logs. As an added complication, some of these logs may be stored on computers remotely located from the system administrator. In addition, format of the various logs may differ from one another, as well as the platforms on which the logs are stored.

A utility available on UNIX systems is Syslog which allows multiple dispersed components to log to a single system. However since it is UNIX only, Syslog does not permit systems having operating systems other than UNIX to write to a common log.

Thus, there is a need, in environments made up of multiple components which operate semi-autonomously, to have the log information generated by these components collected into a centrally located log which can be easily accessed by the system administrator.

SUMMARY OF THE INVENTION

The present patent document discloses techniques for aggregating log data in a distributed system. Previous methods for storing log data have either relied upon
5 maintaining individual logs for each individual process on the local system or a central log for distributed systems wherein each individual system is executed by the same type operating system.

Disclosed in various embodiments are apparatus and methods for gathering event data by a process executed by an operating system on a computer system, transferring
10 that data to a logging process executed by an operating system on another computer system wherein the logging process operating system is intrinsically different from the operating system of the process that detected the event, and storing that data in a data log on the logging process computer system. Provision is also made for gathering, transferring, and storing event data for processes running on the computer system on
15 which the data log is located. A representative data structure for the entries in the data log is also disclosed.

The disclosures of the present patent document provide two primary advantages over the prior art: (1) logging of log data across different platforms and (2) aggregation of log data into one location. This aggregation mechanism is designed to allow multiple
20 elements operating on diverse kinds of systems to log to a central system, which itself may be on any kind of system. An administrator can monitor the operation of a distributed system, as for example a distributed management tool, whose components may be distributed across a network and operating on multiple, geographically dispersed computers.

25 Other aspects and advantages of the present invention will become apparent from the following detailed description, taken in conjunction with the accompanying drawings, illustrating by way of example the principles of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings provide visual representations which will be used to more fully describe the invention and can be used by those skilled in the art to better understand it and its inherent advantages. In these drawings, like reference numerals identify corresponding elements and:

Figure 1 is a drawing of a distributed computer system having a centralized data log as described in various representative embodiments of the present patent document.

Figure 2 is a drawing of another distributed computer system having centralized data log as described in various representative embodiments of the present patent document.

Figure 3 is a drawing of an entry for a data structure for the centralized data log as described in various representative embodiments of the present patent document.

Figure 4 is a flow chart of a method for writing to the centralized data log of figure 1 as described in various representative embodiments of the present patent document.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

As shown in the drawings for purposes of illustration, the present patent document relates to a novel method for aggregating log data in a distributed system. Previous methods for storing log data for distributed systems have either relied upon maintaining individual logs for each individual process on its own local system or a central log wherein each individual system which accumulates log data is executed by a member of the same operating system family. Embodiments disclosed herein are not limited by such constraints. In particular, a process accumulating event data for storage in the central log may be running, not only on a remote computer, but also on an operating system which differs significantly from the operating system of the logging process. Other processes which accumulate such event data in the distributed system may be further executed by operating systems of even different families or types. In the following detailed description and in the several figures of the drawings, like elements are identified with like reference numerals.

Figure 1 is a drawing of a distributed computer system 100 having a centralized data log 105 as described in various representative embodiments of the present patent document. In a first preferred embodiment as shown in figure 1, the centralized data log 105, also referred to herein as the data log 105, is stored in a computer memory 110, also referred to herein as a computer readable memory device 110, on a log computer system 115. A log process 120, also referred to herein as a log program 120, executed by a log operating system 125 stores data in the data log 105. A first computer process 130, also referred to herein as a first computer program 130, is executed by a first operating system 135 on a first computer system 140. The log operating system 125 may differ intrinsically in type from the first operating system 135. When the first computer process 130 detects a first event 145 not shown in figure 1, the first computer process 130 transmits description of the first event 145 as a first event description 150 to the log process 120 via a network 190. However, it is possible that means other than the network 190 could be used to transmit the first event description 150 to the log process 120, as for example storing data on a magnetic disk and physically transferring the disk to the log

computer system 115. The log process 120 stores the first event description 150 in the data log 105. In the first preferred embodiment, the log operating system 125 is intrinsically different from the first operating system 135. It is also possible, that the computer memory 110 comprising the data log 105 could be physically located on a computer system located remotely from the log computer system 115.

Also shown in figure 1 is a second computer process 160, also referred to herein as a second computer program 160, executed by a second operating system 165 on a second computer system 170. The log operating system 125 may or may not intrinsically differ in type from the second operating system 165, and the second operating system 165 may or may not intrinsically differ in type from the first operating system 135. When the second computer process 160 detects a second event 175 not shown in figure 1, the second computer process 160 transmits description of the second event 175 as a second event description 180 to the log process 120 via the network 190. However, it is possible that means other than the network 190, as for example storing data on a magnetic disk and physically transferring the disk to the log computer system 115, could be used to transmit the second event description 180 to the log process 120. The log process 120 stores the second event description 180 in the data log 105. In a representative embodiment, the second operating system 165 is intrinsically different from the first operating system 135. In another representative embodiment, the second operating system 165 is intrinsically different from the log operating system 125. And in yet another representative embodiment, the second operating system 165 is intrinsically different from the first operating system 135 and intrinsically different from the log operating system 125.

Figure 2 is a drawing of another distributed computer system 100 having centralized data log 105 as described in various representative embodiments of the present patent document. In a second preferred embodiment as shown in figure 2, the data log 105 is stored in the computer memory 110 on the log computer system 115. The log process 120 executed by log operating system 125 stores data in the data log 105. The first computer process 130 is executed by the first operating system 135 on the first computer system 140. When the first computer process 130 detects the first event 145

not shown in figure 2, the first computer process 130 transmits description of the first event 145 as the first event description 150 to the log process 120 via the network 190. However, it is possible that means other than the network 190, as for example storing data on a magnetic disk and physically transferring the data to the log computer system 115, could be used to transmit the first event description 150 to the log process 120. The log process 120 stores the first event description 150 in the data log 105. In the second preferred embodiment, the log operating system 125 is intrinsically different from the first operating system 135.

Also shown in figure 2 is an additional log system process 260 executed by the log operating system 125 on the log computer system 115. When the additional log process 260 detects an additional event 275 not shown in figure 2, the additional log system process 260 transmits description of the additional event 275 as an additional event description 280 to the log process 120. The log process 120 stores the additional event description 280 in the data log 105. In a representative embodiment, the additional log system process 260 transmits the additional event description 280 to the log process 120 via the network 190. In another representative embodiment, the additional log system process 265 transmits the additional event description 280 to the log process 120 via paths internal to the log computer system 115.

Figure 3 is a drawing of an entry for a data structure 300 for the centralized data log 105 as described in various representative embodiments of the present patent document. The entry for the data structure 300 comprises a system identification 310 and a component identification 315. The component identification 315 identifies the component which detected the event logged, and the system identification 310 identifies the system on which that component is located. The data structure 300 further comprises event time 320 which specifies the clock time at which the event occurred, and event data 330 which provides information to the log user regarding the nature of the event detected and subsequently recorded in the data log 105. Other items could be included in the data structure 300, as for example operating system and computer system identification. Also, the event time 320 could include the date of the event, as well as the time of day at which the event occurred. Data structure 300 entries into the centralized data log 105 could be

entered into the centralized data log 105 in event time 320 order or as received by the centralized data log 105. They could further be grouped by component identification 315 and/or system identification 310. In practice, the component could be, for example, a software agent, and the system could be the physical system hardware on which the software agent is operating.

Figure 4 is a flow chart of a method 400 for writing to the centralized data log 105 of figure 1 as described in various representative embodiments of the present patent document. The method 400 of figure 4 could be implemented as software processes on distributed computer system 100.

In block 410 the first computer process 130 executed by the first operating system 135 detects the first event 145. Block 410 then transfers control to block 420.

In block 420 the first computer process 130 prepares the first event description 150. Block 420 then transfers control to block 430.

In block 430 the first event description 150 is transmitted to the log process 120 executed on the log computer system 115 by the log operating system 125. Block 430 then transfers control to block 440.

In block 440 the log process 120 receives the first event description 150 from the first computer process 130. Block 440 then transfers control to block 450.

In block 450 the log process 120 stores the event information in the data log 105.

Block 450 is the terminating step in the method.

While the method 400 of figure 4 has been described in terms of the first computer process 130 executed by the first operating system 135 on the first computer system 140, it will be understood that the identical method can be followed for the second computer process 160 of figure 1 executed by the second operating system 165 on the second computer system 170, as well as for the additional log system process 260 of figure 2 executed by the log operating system 125 on the log computer system 115.

In representative embodiments the present patent document describes methods wherein data can be logged across systems of diverse implementations to a log on any kind of system. Thus, the implementations provide two primary advantages over the prior art: (1) logging of log data across different computer platforms and (2) aggregation

of log data into one location. This aggregation mechanism is designed to allow multiple elements operating on diverse kinds of systems to log to a central system, which itself may be on any kind of system. An administrator can monitor the operation of a distributed system, as for example a distributed management tool, whose components
5 may be distributed across a network and operating on multiple, geographically dispersed computers.

While the present invention has been described in detail in relation to preferred embodiments thereof, the described embodiments have been presented by way of example and not by way of limitation. It will be understood by those skilled in the art
10 that various changes may be made in the form and details of the described embodiments resulting in equivalent embodiments that remain within the scope of the appended claims.

What is claimed is:

1. A computer program storage medium readable by a computer, tangibly
embodying a computer program of instructions executable by the
computer to perform method steps for storing event data [330] in a
centralized data log [105] in a distributed computer system [100], the
steps comprising:

detecting a first event [145] by a first computer process [130],
wherein the first computer process [130] is executable by a first
operating system [135];

preparing a first event description [150], wherein the first event
description [150] describes the first event [145];

transmitting the first event description [150] to a log process
[120], wherein the log process [120] is executable by a log
operating system [125], wherein the log operating system [125]
differs intrinsically in type from the first operating system [135];

receiving the first event description [150] by the log process
[120]; and

storing the first event description [150] in the centralized data log
[105].

2. The computer program storage medium as recited in claim 1, providing the first event description [150] is transmitted from the first computer process [130] to the log process [120] via a network [190].
3. The computer program storage medium as recited in claim 1, the steps further comprising:
 - detecting a second event [175] by a second computer process [160], wherein the second computer process [160] is executable by a second operating system [165];
 - preparing a second event description [180], wherein the second event description [180] describes the second event [175];
 - transmitting the second event description [180] to the log process [120];
 - receiving the second event description [180] by the log process [120]; and
 - storing the second event description [180] in the centralized data log [105].
4. The computer program storage medium as recited in claim 3, providing the second event description [180] is transmitted from the second computer process [160] to the log process [120] via a network [190].
5. The computer program storage medium as recited in claim 3, providing the second operating system [165] differs intrinsically in type from the

first operating system [135].

2 6. The computer program storage medium as recited in claim 3, providing the second operating system [165] differs intrinsically in type from the log operating system [125].

2 7. A computer readable memory device [110] encoded with a data structure [300] for transferring data between a first computer process [130] and a log process [120], the first computer process [130] having functions for
4 transferring an event description [150] to the log process [120], the functions having associated parameters, the data structure [300] having
6 entries, each entry containing:

8 event data [330], wherein the event data [330] describes a detected event [145];

10 a component identification [315], wherein the component identification [315] identifies a component detecting the event [145] and wherein the
12 event [145] is described by the event description [150]; and

14 a system identification [310], wherein the system identification [310] identifies a system, wherein the system comprises the component
16 detecting the event [145].

2 8. The computer readable memory device [110] as recited in claim 7, providing the data structure [300] further contains an event time [320], wherein the event time [320] is the clock time of event [145] occurrence.

2 9. A distributed computer system [100] for storing data, comprising:

4 a first computer process [130] executable by a first operating system
[135];

6 a log process [120] executable by a log operating system [125], wherein
the log operating system [125] differs intrinsically from the first operating
8 system [135]; and

10 a centralized data log [105] stored in a computer memory [110], wherein
the first computer process [130] comprises functions for transmitting a
12 first event description [150] to the log process [120] and wherein the log
process [120] comprises functions for receiving the first event description
14 [150] from the first computer process [130] and for storing the first event
description [150] in the data log [105].

- 2 10. The distributed computer system [100] as recited in claim 9, wherein the
first event description [150] is transmitted from the first computer process
[130] to the log process [120] via a network [190].



Application No: GB 0108635.4

Examiner: Michael Powell
Waters

Claims searched: 1 to 10

Date of search: 14 December 2001

Patents Act 1977 Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK Cl (Ed.S): G4A (AFMP, AFMD)

Int Cl (Ed.7): G06F (11/30, 11/32, 11/34)

Other: WPI, PAJ, EPODOC, INSPEC, IBM TDB

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
X	EP 0457110 A2 (IBM) see column 11 lines 1 to 10	1 to 10
X,P	WO 01/20456 A1 (HITACHI) see EPODOC and WPI abstracts	1 to 10
X,P	WO 00/47003 A1 (MPATH INTERACTIVE)	1 to 10

X Document indicating lack of novelty or inventive step
Y Document indicating lack of inventive step if combined with one or more other documents of same category.
& Member of the same patent family

A Document indicating technological background and/or state of the art.
P Document published on or after the declared priority date but before the filing date of this invention.
E Patent document published on or after, but with priority date earlier than, the filing date of this application.